

Belle II Software

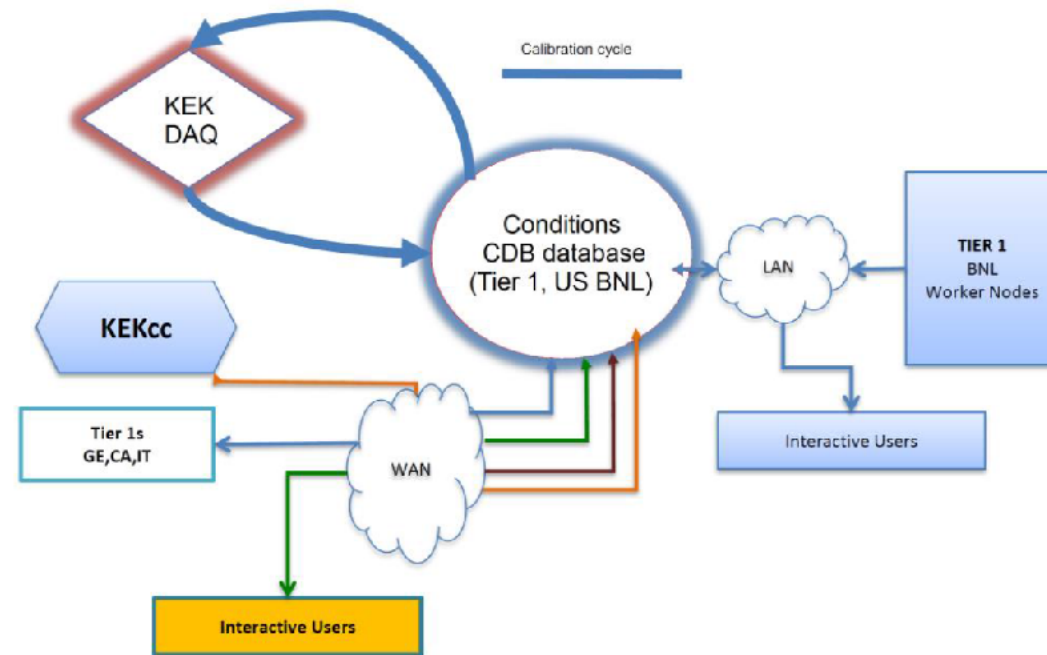
Paul Laycock for the Belle II team

**(The work shown is done by the whole Belle II team,
thoughts impressions and musings are my own!)**

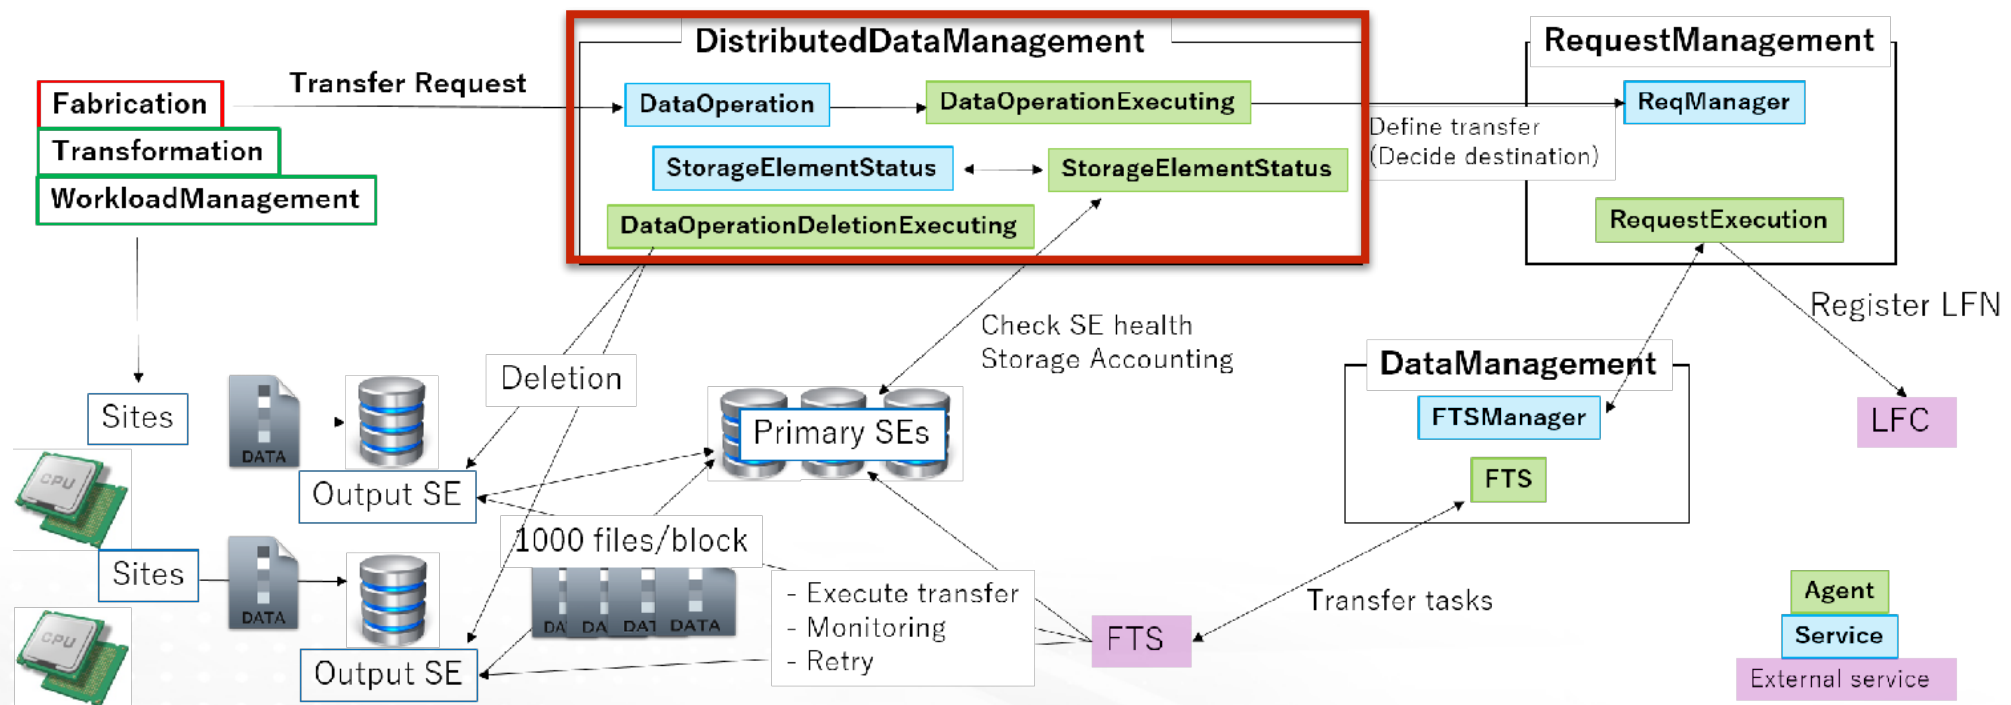
BROOKHAVEN
NATIONAL LABORATORY



Conditions database (CDB)

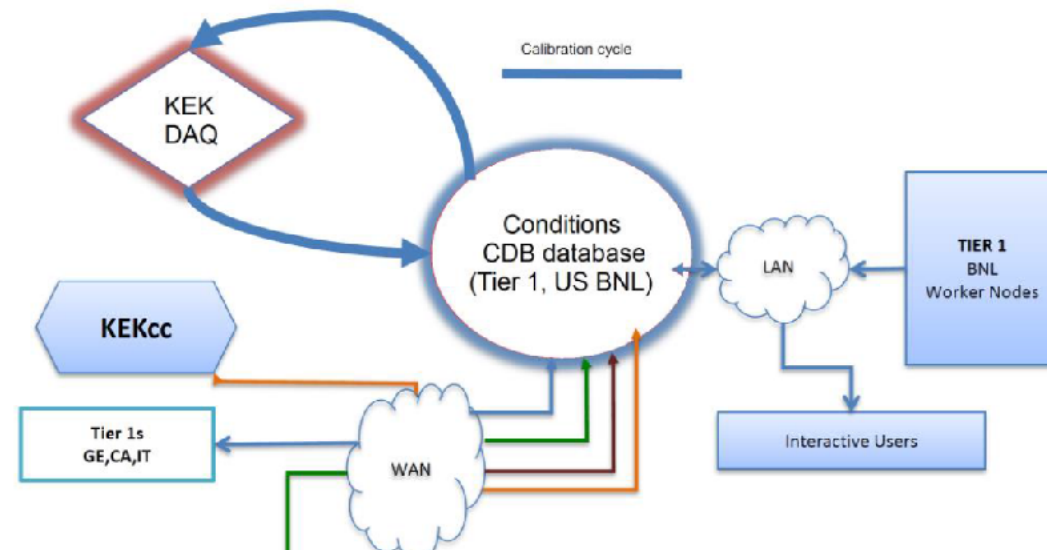


Distributed data management (DDM)



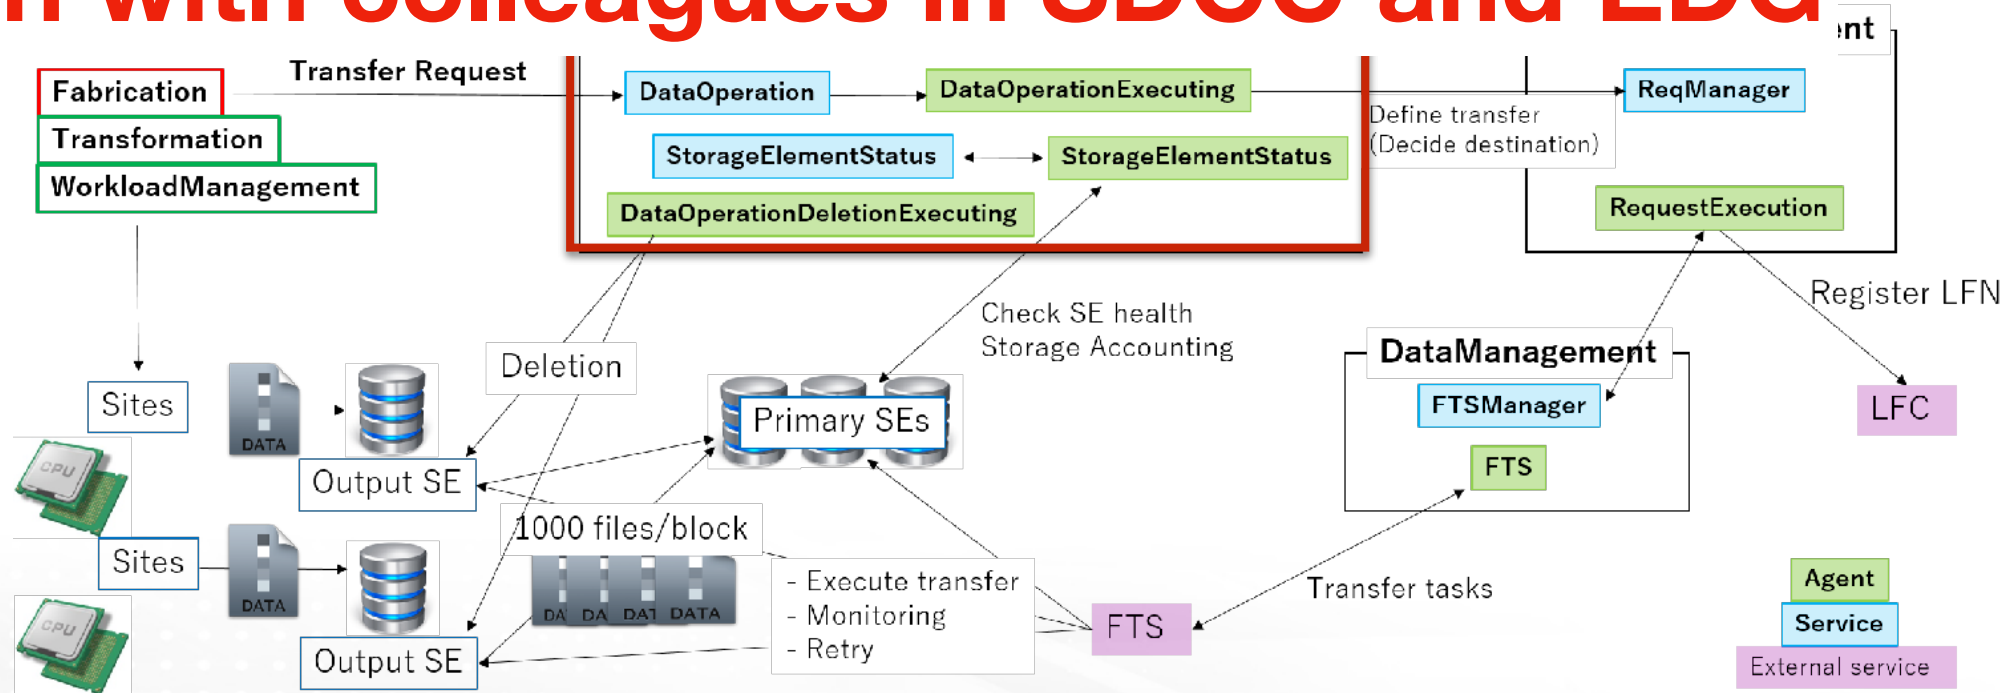
BNL Belle II software responsibilities

Conditions database (CDB)



Of course, all of this work relies on close collaboration with colleagues in SDCC and EDG

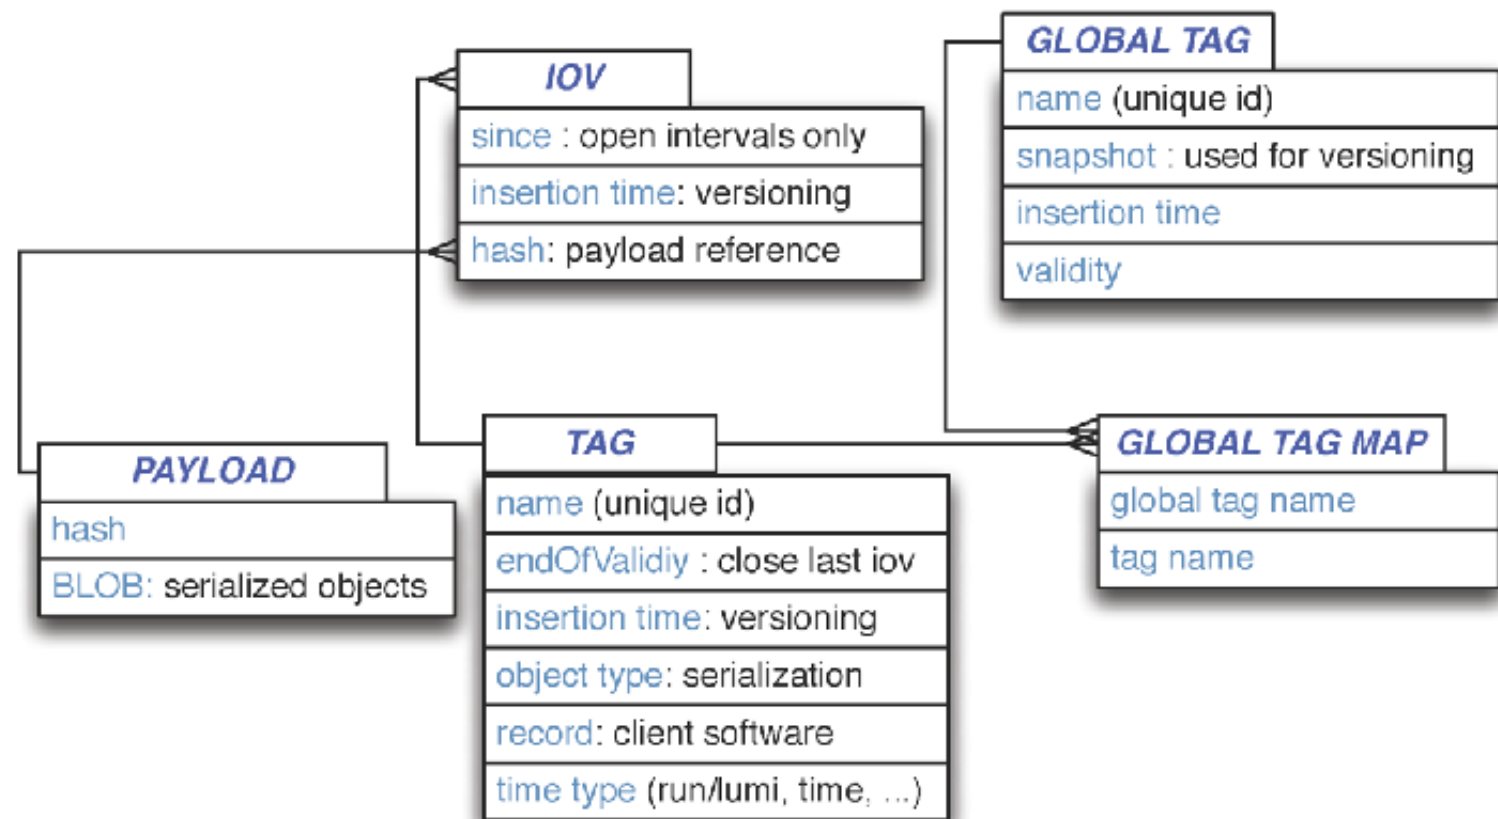
Distributed data management (DDM)



Belle II Conditions Model

- Data Model: relational DB
- Single tables for payload, tags, IOVs
 - ***Payloads can be separated completely from metadata***
 - ***Largely experiment agonistic***
- *IOVs and payloads resolved independently*
- **Cache-friendly design**
- Largely follows best practice principles in **HSF CWP** paper:

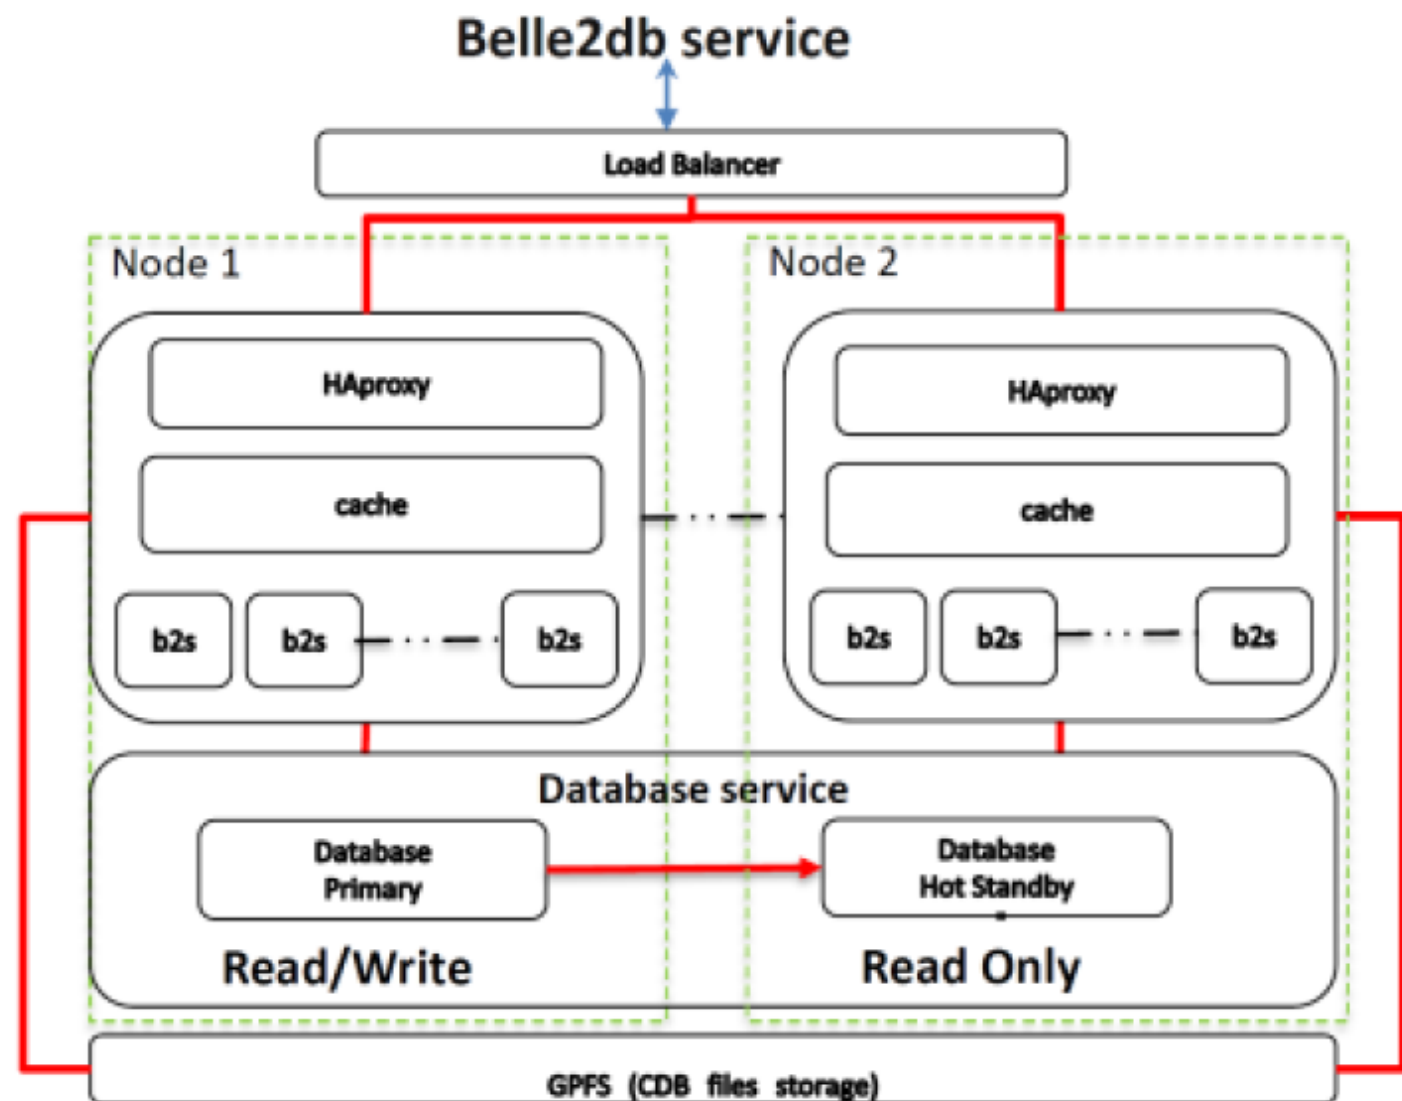
- <https://arxiv.org/abs/1901.05429>



Data model from CWP paper

Belle II Conditions Service

- Containerised component architecture
 - *Highly scalable*
- Industry-standard components
 - **REST** interfaces
 - *Swagger auto-generates API*
- ***Payloads entirely factorised***
 - Postgres for metadata, GPFS for payloads
 - Separate file service for reading
 - Payloads are also cached locally and on cvmfs
- Good performance, stress-tested ***beyond expected BelleII requirements***
- DB schema slightly different but principle is the same
 - *largely experiment agnostic*



Belle II CDB possibilities

- CDB is in good shape, but some key things were missing, especially ***authentication and authorisation*** - work in progress
- Workflows and Belle II specific needs will likely require more work through 2019 as experience with calibrations improves - physics data-taking aka *phase 3* only just started in **March 2019**

Nevertheless

Belle II CDB could be a blueprint for a generic CDB software

- Making the code ***experiment independent*** could open the possibility of making this CDB design available to ***more experiments***
- Would also require work to define ***experiment-dependent workflows and end-user tools***
- Use cases - this is particularly appropriate for ***high rate access*** patterns of offline distributed computing, where ***caching is critical***

Belle II Distributed Computing

Definition

- MC prod / data process
- Type (BB, $\tau\tau$, ccbar..)
- # of events
- software version
- etc..



PS



-Production
-Distribution
-Merge

Production manager (human)
- Define "Production"

Production Management

Belle
DIRAC

Distributed data management

- Gather outputs to major storage (and distribute over the world)
- Check status of storages
- Define "Transfers"

DIRAC

Request management

Data management



Primary SE

Resource

B2Monitoring

output info

Fabrication

Transformation

- Define jobs
- Re-define failed job
- Verify output files


Workload Management

- Submit job on site
- Record job status etc

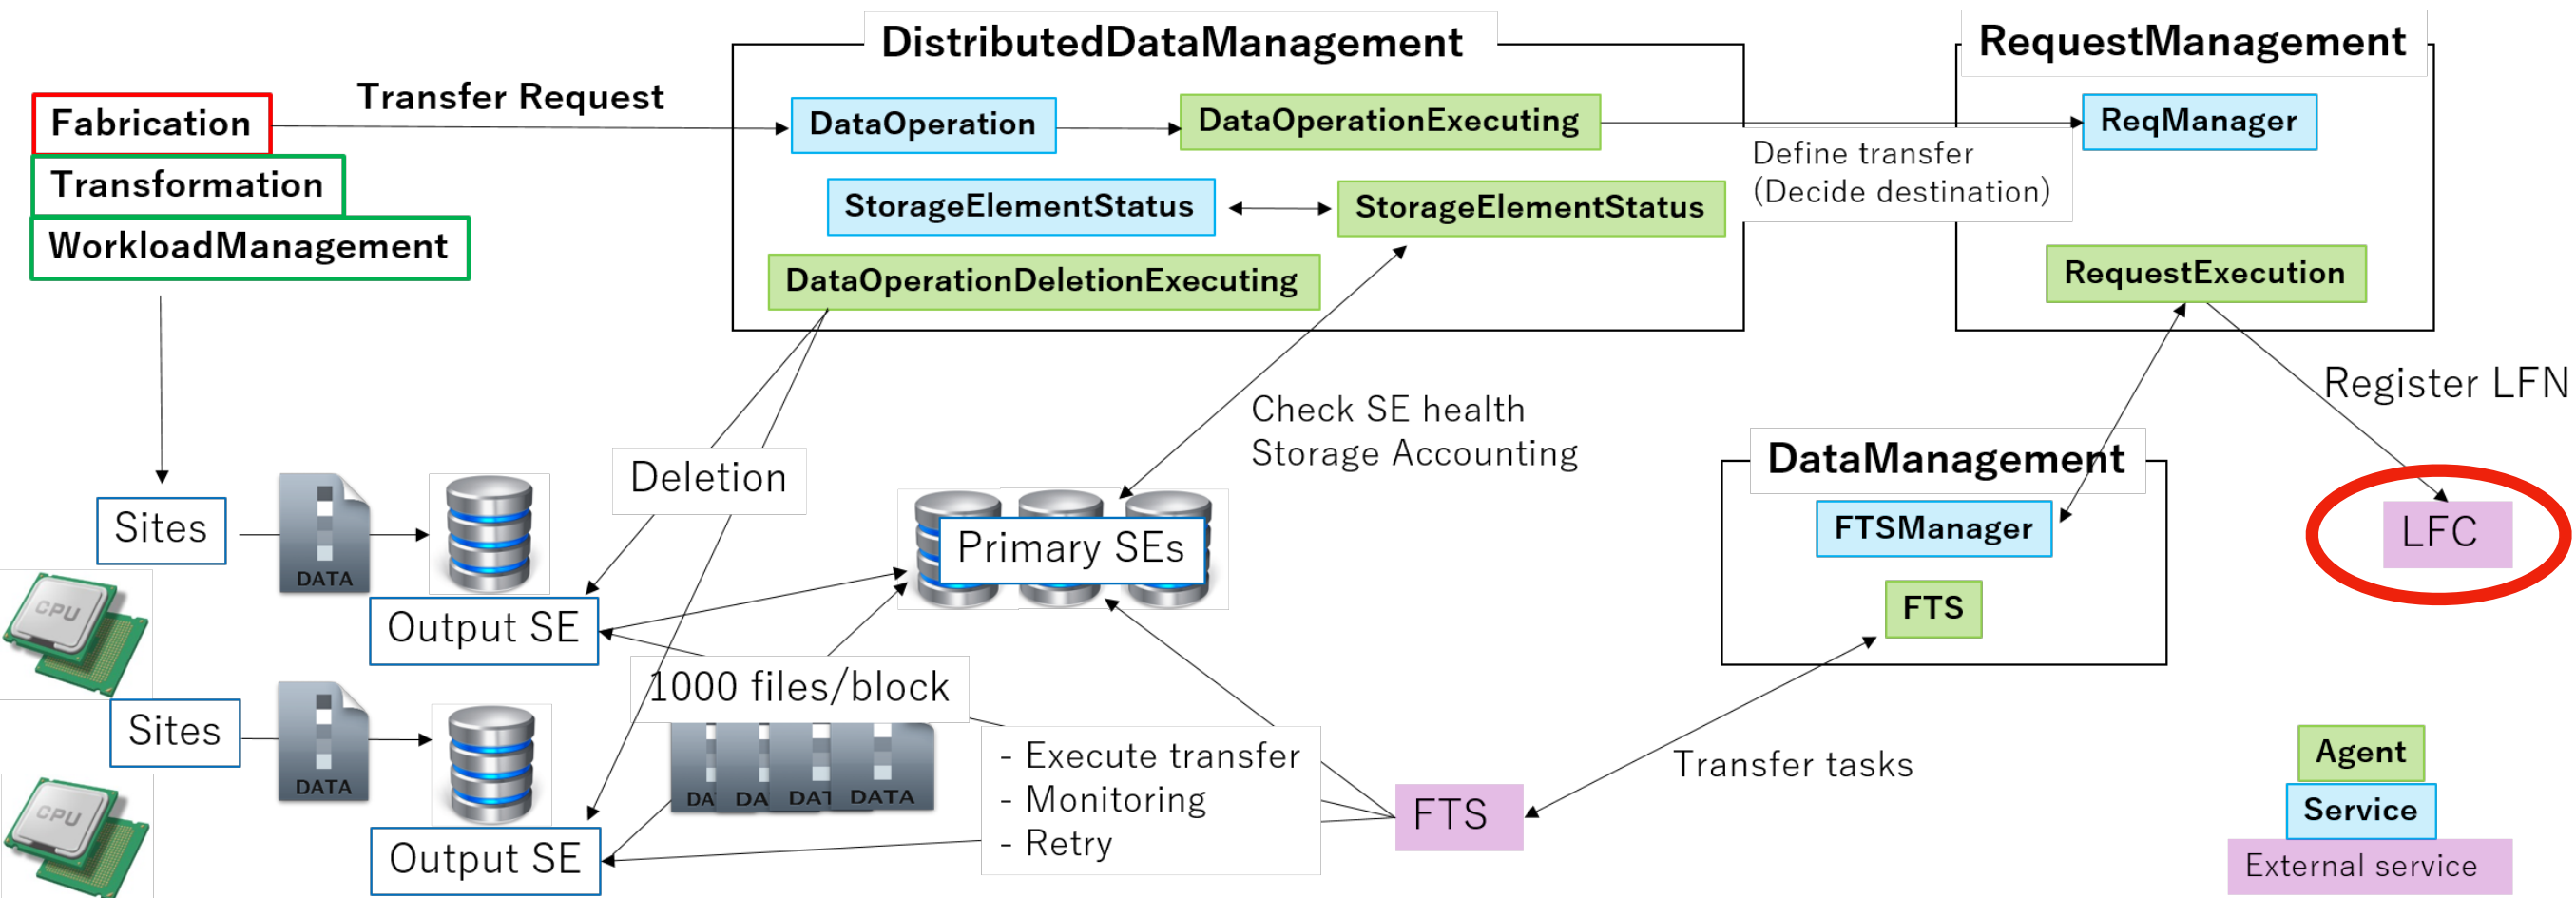
Output SE



Computing site

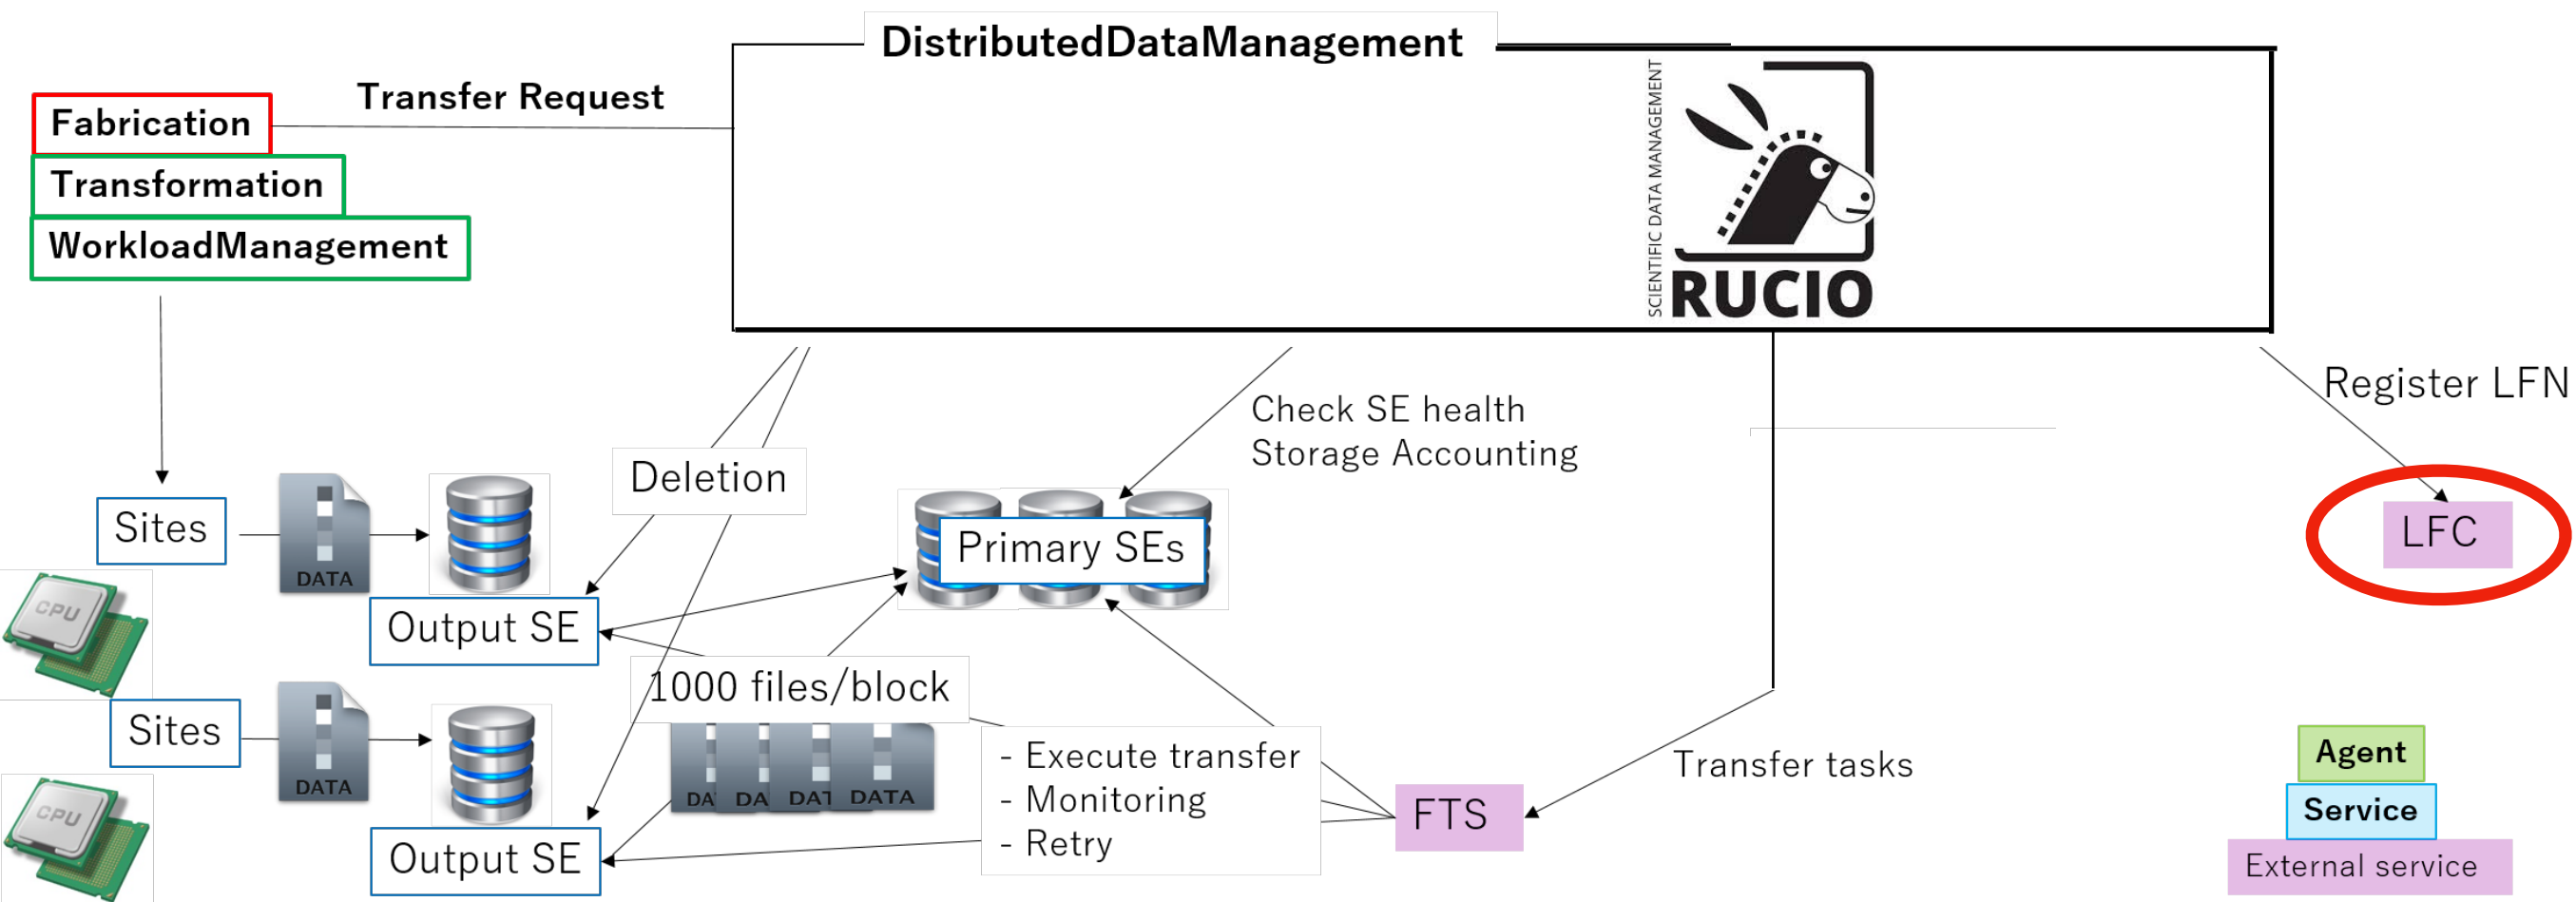


Belle II Distributed Data Management



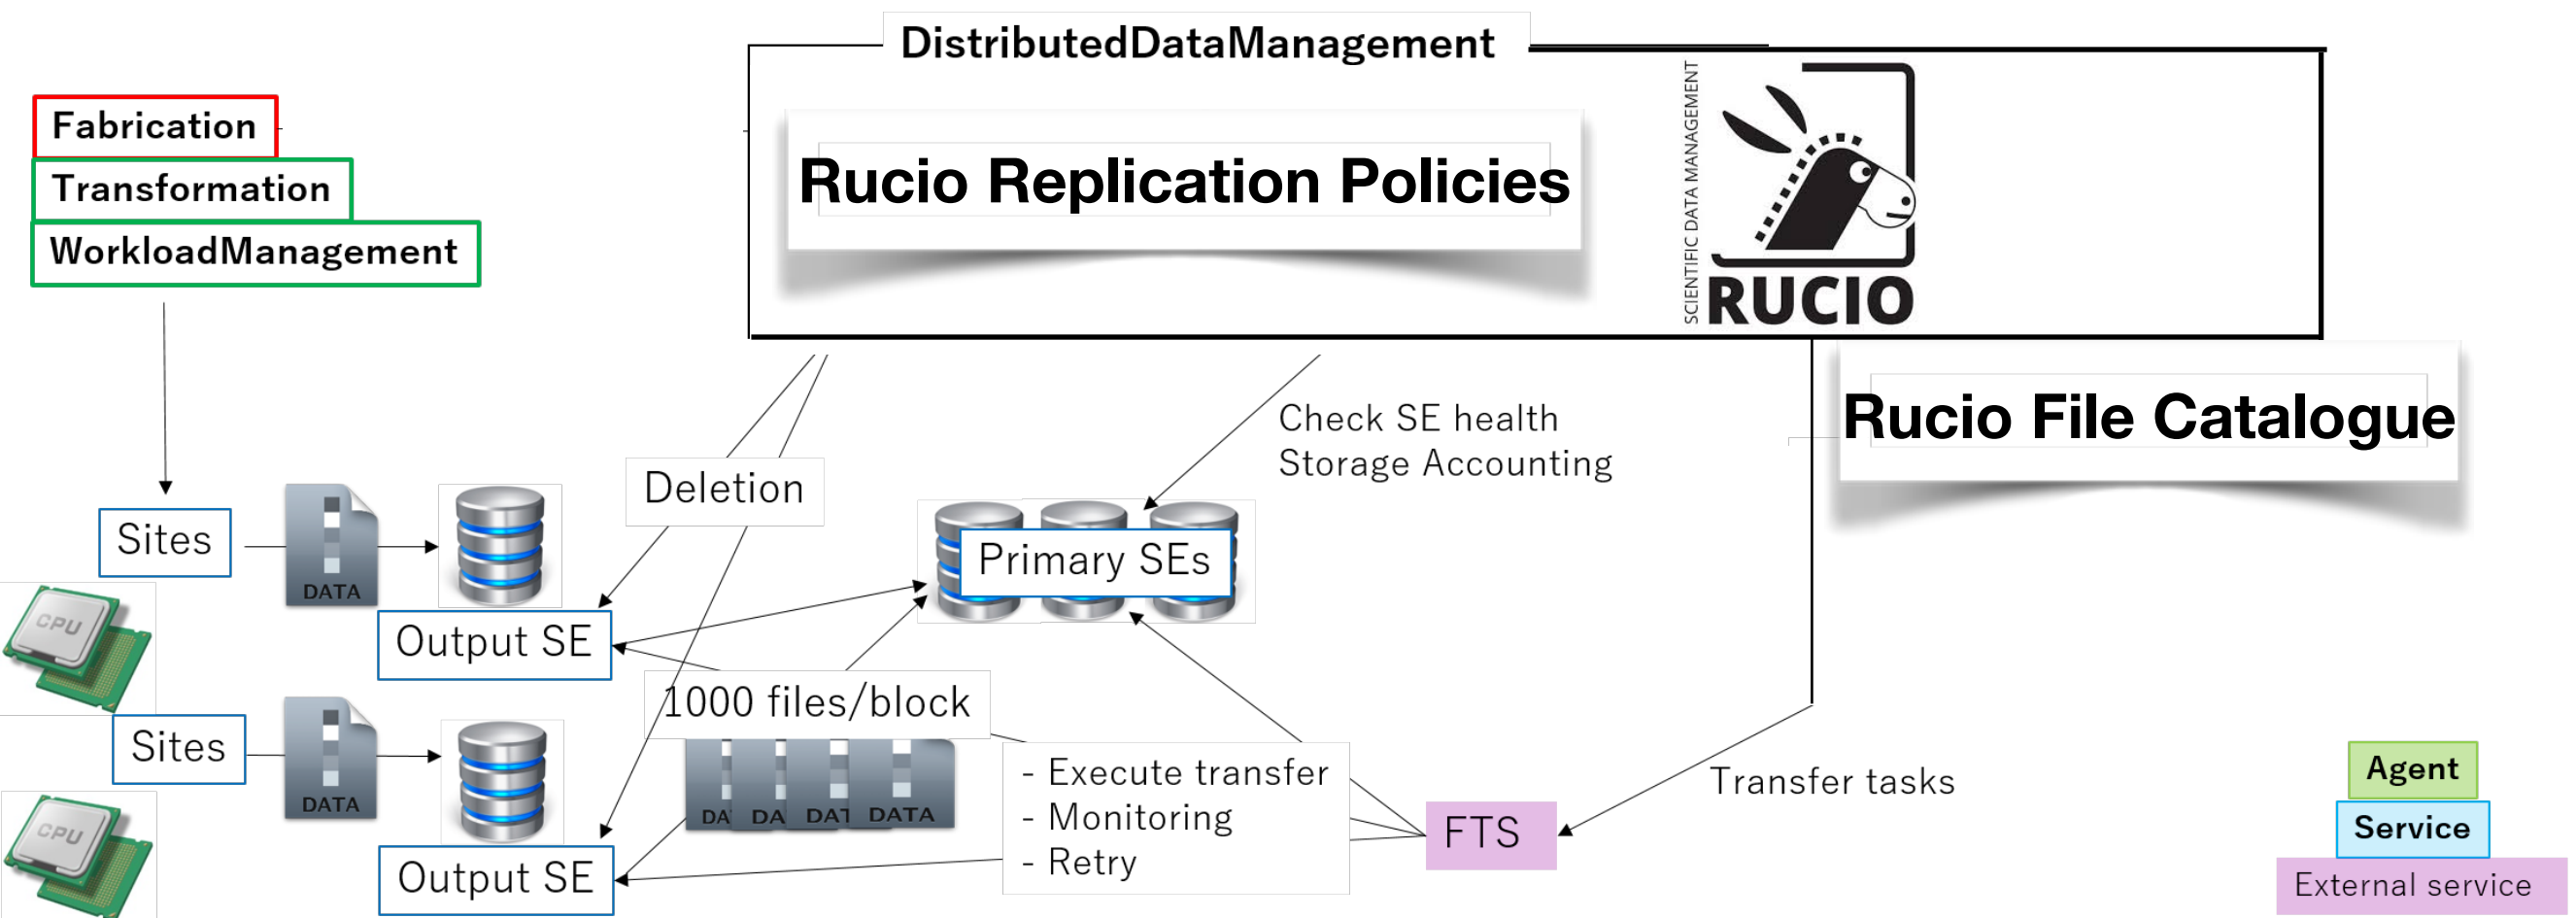
- Designed as a key component of **BelleDIRAC**, the **DDM** is well integrated into the **DIRAC** ecosystem (production system, monitoring, et al)
- However, only basic functionality was mature (and still required work), key features either untested or missing, and the **LFC file catalogue** is soon to be extinct

Belle II Distributed Data Management - Future I



- Strong desire to move **Rucio as DDM**
- Evaluation: no show-stoppers but **priority was to fix and operate current DDM**
- In first stage migration, ***maintain current API*** to minimise impact
- **Rucio** is used *under the hood* but **LFC** is still *master file catalogue*

Belle II Distributed Data Management - Future II



- Second stage migration, move to use **Rucio as the master file catalogue**, interfaced to DIRAC using **file catalogue plugin**
- Rely on replication policies to define **Rucio rules**, **lifetime policies** to **automate deletion**
- Remove explicit calls to **DDM**
- Use of **file catalogue plugin** should **avoid most work on user side**, but may be able to make more use of **Rucio** by **improving end-user tools**

Belle II DDM possibilities

- Current **BelleDIRAC DDM** performance is ok, but many features missing, and large operations overhead due to lack of automation
- Window of opportunity to “*just use Rucio*” was very small and passed, too high risk at the start of data-taking, ***rely on current DDM in 2019***, so for now we have operations burden and development effort can only really shift to **Rucio** once current **DDM** is good enough
- *Looking ahead, move to **Rucio** will reap rewards*
 - Some tooling needed to make full use of advanced monitoring, use lifetime policy, etc. and detailed understanding of **Rucio** operations will take time
 - ***Development effort should largely be under control, no reinventing the wheel***
- Integration with **DIRAC** very popular in the field, several experiments and communities mentioned this at ***Rucio workshop***, including **DUNE** and **SKA**
- Hot topic for ***DIRAC Users workshop next week***
- **BNL** could be unique in having **PanDA-Rucio** and **DIRAC-Rucio** expertise
- Good to stay agile, both of those combinations have their advantages, and both have a common denominator - ***Rucio expertise will be needed for the coming decade***

Expectations and Conclusions

- Significant contributions from **BNL** to **Belle II** software
 - *Thanks go to the **whole** team at **BNL***
- **Expectations**
 - We are fully committed with the current level of effort, 2.5-3 FTE in NPPS
 - The creation of NPPS does not change this today, but optimistic about the future
- **Opportunities**
 - **CDB** is in good shape, if effort existed, potential to take this as a blueprint and offer to other interested experiments, those hosted here at **BNL** and e.g. **DUNE**
 - **DDM** is more challenging, but the challenges here also represent opportunities
 - **DIRAC-Rucio** integration is going to happen, opportunity for **BNL** to play a leading role in this
 - ***Rucio has established itself as the DDM tool, Rucio expertise is a valuable resource***
 - Further ahead, **Belle II** is something of a poster child of analysis software:
 - ***Declarative analysis***, using ***jupyter notebooks*** at ***analysis facilities***
 - ***Potentially*** opportunities to contribute to data analysis tools together with **SDCC**
- *Looking forward to seeing how things evolve in the near future*